

**IN THE SPECIFICATION:**

Please amend the specification as follows:

Please insert the following figure description after the description for FIG. 9 on page 9.

FIG. 10 shows selected frames of classification results.

Please replace the last paragraph on page 18 with the following amended paragraph.

Second, the RTWCS's analyzer is structured as a pipeline of six analysis modules (see FIG. [[4]] 4A). If a cell image does not meet the condition in a particular module, it will not be further processed by the subsequent modules in the pipeline. Therefore, many cell images do not go through all six modules. At the end, only a small number of cell images containing the newly appeared pen strokes come out of the analyzer. The six modules are:

1. Change detector 402: This module determines if the cell images have changed since last frame.
2. Color estimator 404: This module computes the background color of the cell images -- the color of blank whiteboard.
3. Background modeler 406: This is a dynamic module that updates the whiteboard background model by integrating the results computed from the previous modules which may have missing parts due to occlusion by foreground objects.
4. Cell classifier 408: This module classifies the cell images into foreground or whiteboard cells.
5. Stroke extractor 410: This module extracts the newly appeared strokes.
6. Color enhancer 412: The color enhancer enhances the color of the newly appeared extracted strokes.

The change detector, color estimator, background modeler and cell classifier modules all preferably operate one image data in Bayer color space.

Please replace the second paragraph on page 25 with the following amended paragraph.

Note that one increases the uncertainty of its neighbors by  $\Delta$  (4 in our system) to allow color variation. A hole of size  $N$  generally takes  $N/2$  frames to get filled. Since the uncertainty in the cells with filled values is much larger than the ones with the observed values (due to added  $\Delta$ ), the filled values are quickly supplanted by the observed values once they become available. An example of an integrated whiteboard color is the third image in FIG. [[8]] 8C. Note that the bookshelf area in the left side of the image is never filled.

Please replace the third paragraph on page 26 to with the following amended paragraph.

However, more accurate results can be achieved by utilizing spatial relationship among the cell groups. The basic observation is that foreground cells should not appear isolated spatially since a person usually blocks a continuous region of the whiteboard. So at the neighborhood level, the RTWCS performs two filtering operations on every frame. First, the RTWCS identifies isolated foreground cells and reclassifies them as whiteboard cells (process actions 908, 910). This operation corrects the misclassification of the cells that are entirely filled with strokes. Second, the RTWCS reclassifies whiteboard cells which are immediately connected to some foreground cells as foreground cells (process action 912). One main purpose of the second operation is to handle the cells at the boundaries of the foreground object. Notice that if such a cell contains strokes, the second operation would incorrectly classify this cell as a foreground object. It will be correctly re-classified as whiteboard once the foreground object moves away. Extending the foreground object boundary delays the recognition of strokes by a few frames, but it prevents some parts of the foreground object from being classified as strokes--a far worse situation. FIG. 10 provides samples of classification results.